

METHODS AND SYSTEM FOR DATA PACKET FILTERING USING TREE-LIKE HIERARCHY

TECHNICAL FIELD

5 **[001]** The invention pertains to packet filtering. More specifically, the invention relates to a method of data packet filtering that utilizes a decision graph having a tree-like structure to allow fast and efficient packet filtering.

BACKGROUND OF THE INVENTION

10 **[002]** Internet protocol (IP) networks have been deployed at an increasing rate as a transport mechanism for a large number of data networking applications. IP packet filtering is a process of checking each IP packet that is going to be sent from one computer to another computer in an IP network. Essentially, a network consists of nodes, and links coupling the nodes to each other, with suitable data switching, routing, reception and 15 transmission equipment at the nodes. IP packet filtering is widely used in internet firewall systems by independent service providers and organizations connected to the network.

15 **[003]** Filter rules are most commonly employed to allow implementation of differentiated services based on the data type. One problem that must be overcome for a differentiated 20 services implementation is packet classification. That is, for a given incoming packet, a selection as to which of several service schemes are to be applied to the packet must be made by selecting one filter among many which the packet may match. For packet classification, a classification key composed of packet fields located at pre-specified offsets within the packet is typically used as a packet filter. A packet filter may specify 25 conditions for many items in the packet stream such as the IP, TCP, UDP, or ICMP headers, including the IP source address (IPSA), IP destination address (IPDA), Differentiated Services Code Point (DSCP), IP protocol field, source port, destination port, and IP precedence field. Each of the subfields of the IP filters are specified by giving a

value and a mask length, specified in the number of bits from the most significant bit to be considered and the value it must have for this filter to match.

5 [004] It is important and necessary to have efficient filtering algorithms which can determine quickly which of the filter conditions match the current arriving packet. Speed is of the upmost importance as the decision must be made for every packet that arrives at the equipment interface. Typically, the equipment comprises a router with an interface to one or more data links of a data network. The router bears the primary responsibility for inspecting the incoming data packet and applying filtering algorithm that enables 10¹⁴ differentiated services. At the same time, there are many other uses of a filtering algorithm such as, for example, determining which packets belong to certain tunnels in the IP architecture, access filtering, policy based routing, and accounting filters. In addition, there are non-networking applications which may utilize packet filtering algorithms for solving a wide variety of regular expression matching problems.

15 [005] Thus a need exists for an improved method of filtering incoming data packets. An efficient and fast way of filtering IP packets for packet classification allowing the implementation of differentiated services would provide numerous advantages.

20 SUMMARY OF THE INVENTION

25 [006] The present invention provides a filtering mechanism that can be applied to an incoming data packet stream to achieve fast and efficient packet filtering and thereby enable the implementation of differentiated services and other expression matching applications. The invention utilizes a decision graph having a tree-like hierarchy to narrow down a set of matching filters at each stage of the graph. With the present invention, the number of comparisons needed to determine the subset of matching filters which match an incoming data packet stream is equal to the number of bytes in the search and not the

number of filters. This provides a filtering solution that is scalable and efficient in terms of decision processing time.

[007] Therefore, according to one embodiment, disclosed is a method of filtering an incoming data packet in a data packet transmission system. The method comprises a step of creating a decision tree having a tree-like hierarchy organized into edges and stages having test conditions which define filters. The filters are arranged along edges and stages of the tree-like hierarchy according to common values. When a data packet is received at a point of ingress to the data packet transmission system, condition values along the edges of the tree-like hierarchy are compared to specified portions of the data packet to determine where there are matches. If matches are made along an entire edge of the tree-like hierarchy, the data packet is filtered according to the number of comparisons made and not the number of filters.

[008] Also disclosed is a method of filtering incoming packets at the interface of a router to provide differentiated services when the router receives incoming packets from a data network. The method comprises the step of constructing a data structure that stores a set of filters in a tree-like hierarchy. The tree-like hierarchy has edges organized in two stages with the set of filters subdivided according to their values along the edges and stages.

Next, the interface of the router receives an incoming data packet and compares segments of the incoming data packet to values of the filters stored in the tree-like hierarchy. The comparison is done such that data packets can be differentiated by comparing specific packet fields to the set of filters which have been organized according to common values along the tree-like hierarchy meaning segments are compared along stages, and not every edge of the tree-like hierarchy, thereby allowing the incoming data packet to be differentiated according to matches between segments compared and specific filter values. The comparison step can be done by executing a single instruction that filters incoming

data packets according to their contents. Also, the comparison step can be done on a byte-by-byte basis with an offset from the beginning of the incoming data packet to increase the speed of filtering. Once a match is made at one byte, a second byte of the data packet is compared at a lower stage of the tree-like hierarchy, thereby eliminating the
5 need to perform a comparison against all possible sets of filter conditions.

10 [009] Also disclosed is a router for use in a data network that allows the implementation of differentiated services by efficient filtering of incoming data packets. The router comprises an interface to the data network and a packet processing unit coupled to the
15 interface and adapted for receiving incoming data packets via the interface from the data network. The packet processing unit may comprise an application specific integrated circuit (ASIC) designed for filtering purposes according to the kind of services and level of differentiation required by the system. The router further includes memory means operably coupled to the packet processing unit and adapted for storing software instructions that cause the packet processing unit to filter incoming data packets. A data structure is provided and adapted for storing a set of filters in a tree-like hierarchy, the tree-like hierarchy including edges and stages, with the set of filters subdivided according to their values along the edges and stages of the tree-like hierarchy. The software instructions in the memory space are adapted to cause the packet processing unit to
20 compare segments of incoming data packets to values of the filters stored in the tree-like hierarchy, such that the segments are compared along stages and not every edge of the tree-like hierarchy. This allows the incoming data packet to be differentiated according to matches between segments compared and specific values as opposed to all possible filters.
25

[010] The present invention offers numerous advantages over conventional ways of packet filtering. With the present invention, the filtering algorithm is scalable and the

number of instructions executed to filter a packet does not increase as the number of filters grows. Moreover, the methodology of the present invention can be implemented in an algorithm that uses a single software instruction with parameters that greatly aid in the speed and efficiency of the filtering function.

5

BRIEF DESCRIPTION OF THE DRAWINGS

[011] The above, and other advantages, will be understood by reference to the following detailed description in connection with the appended drawings in which:

10 **[012]** Figure 1 is a process flow diagram for the method of the present invention of filtering incoming data packets;

[013] Figure 2 illustrates the use of filter sets and their organization along a tree-like hierarchy;

15 **[014]** Figure 3 illustrates the tree-like hierarchy and having filter sets with wildcards;

[015] Figure 4 illustrates the variables that can be used with a single instruction for packet filtering;

20 **[016]** Figure 5 is a process flow diagram for the method of grouping filters into sets and byte subdivision according to the present invention;

25 **[017]** Figure 6 is an example illustration of the use of a single instruction to perform packet filtering according to the present invention; and

[018] Figure 7 is a generalized block diagram of a router adapted for performing the packet filtering method of the present invention.

5 **[019]** Corresponding numerals and references in the figures correspond to like references in the detailed description unless otherwise indicated.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

10 **[020]** The present invention provides methods and a related system for filtering an incoming data packet using a data structure that has a tree-like hierarchy. With the present invention, filters are subdivided according to their common values along particular portions of the filters. Subdivision continues until all common values have been organized along the tree-like hierarchy which is used to filter incoming data packets within the data transmission system.

15 **[021]** With reference now to the figures and in particular to Figure 1, therein is shown a process flow diagram of the method of filtering IP packets according to one embodiment of the present invention. The method, denoted generally as 10, provides a way of defining a decision graph which can be used to narrow down a set of matching filters for efficiently filtering incoming data packets. In general, the method 10 operates by limiting the number 20 of decisions to the number of bytes of the filter as opposed to the number of filters.

25 **[022]** Filters consist of a series of value/mask pairs where the value represents some series of bits that are compared against sections of an incoming data packet. The mask specifies which of the bits in the value field are significant. For method 10, all masks are left contiguous in the filter field and a notation can be used to designate the value/mask pair. For instance, if a match is desired for any incoming source address (IPSA) which is four bytes in length with the first two bytes given in its dotted decimal form (such as

192.168), the mask can be specified as 255.255.0.0. For simplicity, the value/mask pair designation can be written as 192.168.0.0/16, representing the fact that a match is desired with the 16 leftmost bits. Given a desired match with 192.168, a first filter can be defined as 192/8 for the first byte of the IPSA and 168/8 for the second byte. The remaining two 5 bytes of the IPSA are wildcarded and would be designated at 0/0 since the value of these two bytes are a “don’t care.” In this way, a filter is defined which can be compared and matched against specific bytes of incoming data packets.

10 [023] Therefore, method 10 begins at step 20, wherein the set of all filters of interest are defined according to the type of packet classification desired. Next, at step 22, the set of filters is divided based on specified values of portions of the filters such as, for example, specific bytes of each filter. Next, at step 24, a decision is made if all filters have been visited and, if not, process flow is redirected back to step 22 wherein the filters are further subdivided according to subsequent portions of the filters, such as by locating another byte 15 position of the filter. This process continues until all filters have been visited.

20 [024] Once all filters have been visited, process flow is directed to step 26 wherein the decision graph structure has been completed. We refer to the decision graph structure as a decision tree, its use to be described in further detail below. Next, at step 28, the system waits to receive an incoming data packet. Once received, portions of the incoming data 25 packet are compared with edges of the decision tree, step 30. Since condition comparisons are made along edges and stages of the decision tree, the number of comparisons made is equal to the number of filter bytes to the number of filters. Specifically, as explained in more detail below, the decision tree is organized into levels or “stages”, such that once a match is made at one stage, the filtering algorithm automatically drops to a second stage without further comparison at the stage above. If matches are made along an entire edge of the decision tree, the data packet is filtered and

its contents have been classified according to value. Therefore, a match has been made, step 32. By filtering an incoming data packet in this way, the data system is able to efficiently classify a packet and provide fast differentiated services.

5 [025] Therefore, the present invention provides a method of filtering a data packet in a data packet transmission system. With reference now to Figure 2, therein is shown the process of the present invention utilizing a tree-like hierarchy that represents a decision tree for filtering incoming data packets. The tree-like hierarchy, denoted generally as 50, is used to divide a set of filters 60 into a plurality of stages 70a:70e starting with an 10 uppermost stage 70a to the lowest stage 70e. Each stage corresponds to a specified byte 72a:72e of the filters 60, such that the filters are subdivided along edges of the tree-like hierarchy 50 in terms of common values. It should be understood that tree-like hierarchy 50 is representative of a decision tree in the most general sense and that specific structure 15 for a decision tree would depend on the system and level of classification required.

15 [026] As shown, filter sets 60 contain a plurality of filters (1, 2, 3, 4, 5, 6, 7) with each filter 20 in turn having specified values at each byte (A, B, C, D, E, F, Q, X, Y, Z). Therefore, the process of subdividing the filter sets 60 to facilitate packet filtering begins by choosing a single byte, for example byte 1, which is subdivided based on its values. The resulting 25 division, as shown in stage 70, leads to filters whose first byte contains an "A" value and a second filter whose first byte contains an "X" value. In this example, edges 74 and 76 define paths of the tree-like hierarchy 50 with common values at the same filter position.

25 [027] While the example of Figure 2 illustrates the subdivision methodology of the present invention using seven filters 60, it should be understood that more or less filters may be utilized in any particular implementation. The method of subdividing and grouping filter sets facilitates a wide array of differentiated services applications as well as an efficient

and fast way of data packet filtering. Likewise, while the tree-like hierarchy 50 includes five stages 70a:70e, it should be readily understood that more or less stages may be employed in any particular packet filtering algorithm.

5 [028] The process of subdividing the filter 60s continues as illustrated in Figure 2 to create the tree-like hierarchy 50 having five stages 70a:70e corresponding to five bytes 72a:72e, such that the lowest stage 70e represents the final possible subdivision of filters having common values. Thus at the lowest stage 70e, all filters have been visited along all relevant byte positions. The tree-like hierarchy 50 represents the building of a decision
10 tree which can be used to quickly determine which subsets of filters match a particular incoming packet. When a particular packet arrives, the algorithm of the present invention can be used to filter the data packet at the position representing byte 1 of the search key to determine whether it contains an "A" or an "X" corresponding to edges 74 and 76, respectively. If the incoming packet contains an "A" expression, matching would continue
15 from edge 74 to edges 78a, 78b or 78c to determine if the incoming data packet at the second byte includes values "B," "Y" or "Q."

20 [029] The process continues in this way walking down all stages 70a:70e down the tree-like hierarchy 50, after which it is known exactly which filters match the values of the incoming data packet. Therefore, the comparison methodology of the present invention provides an efficient and fast way of determining the subset of matching filters since the number of comparisons is equal to the number of bytes that are being compared and not the number of filters. This provides a packet filtering algorithm that is scalable in terms of decision processing time.

25

[030] With reference now to Figure 3, therein is shown an example tree-like hierarchy 100 which is similar to the tree-like hierarchy 50 of Figure 2, except that it accommodates byte

values with wildcards meaning that the filter is not specific to a value at one or more positions of the incoming data packet. This is shown by wildcarded filters 102 wherein the "don't care" values are shown as an asterisk. Note that filter 1 of wildcarded filters 102 has wildcard values at every position, meaning that any incoming data packet is matched by 5 any input that matches none of the other subsets.

10 **[031]** The tree-like hierarchy 100 is more complicated than the tree-like hierarchy 50, resulting in a more complicated decision tree including wildcarded bytes that get replicated across all subsets of a parent filter set. Compared to linear matching expression algorithms, however, the process using a tree-like structure to subdivide filters, as illustrated in Figures 2 and 3, provides significant efficiency advantages as well as speed in the matching algorithms and scalability.

Byte selection heuristic

15 **[032]** It should be apparent to one of ordinary skill in the art that the examples of Figures 2 and 3 are but exemplary and that changes to the structure and form of the decision tree used to enable packet filtering can be made. For example, with the tree-like hierarchies 50, 100, the first byte was chosen to enable the comparison on an incoming data packet in a sequential order. While this simplifies the logic structure of the decision tree, it will 20 rarely lead to an optimal tree in terms of its total tree size. It has been found, therefore, that the size of the decision tree can be reduced by dynamically choosing the next byte for comparison along any edge using a heuristic. Moreover, it has been found that a good heuristic is one where the next byte chosen contains the fewest number of wildcarded filters since the wildcarded filters are replicated in each subset in a parent filter set by 25 choosing a byte containing the fewest number of wildcarded filters. Subsets are created from the filter set containing the smallest number of entries when summed together. Hence, the subtree tends to be smaller in size.

5 [033] Therefore, the present invention provides a way of limiting the size of the tree-like structure used for a decision tree during IP packet filtering by dynamically choosing the best byte at each stage, such that different paths down the tree-like architecture, from the root to a leaf, may compare bytes in a different order.

Collapsing subsets

10 [034] Another aspect of the tree-like architectures 50, 100 is that the decision graphs are formed by successively subdividing a set of filters according to a byte in the search key. For process optimization, however, it is possible to collapse like subsets at any stage in the packet filtering algorithm. Therefore the decision graph actually may form a directed acyclic graph, as multiple edges from a parent may collapse into the same child.

Single instruction

15 [035] The packet filtering methods of the present invention include novel components in that a comparison is done at each stage of the decision tree (as represented by Figures 2 and 3, as examples). The algorithms, however, may be implemented in any suitable device, such as an application specific integrated circuit (ASIC) designed to perform the filtering functions. It has been found that an ASIC so designed can work in connection with 20 a single software instruction which aids in the filtering process. An example of such an instruction is as follows:

`cmpdemux(cmp_len, cmp_val, demux_len, pkt_offset, table_ptr), cmp_fail_ptr`

25 [036] Therefore, as illustrated in Figure 4, the *cmpdemux* instruction advances from the beginning of the packet's IP header field 110 the value specified in the packet offset (*pkt_offset*) field 112. The instruction looks at the next *cmp_len* bits and determines if the

value of the bits in the *cmp_val* field 114 is *cmp_val*. If it is not, execution jumps to the address specified by *cmp_fail_ptr*, meaning that the filtering algorithm has been unsuccessful. Otherwise, the comparison succeeds and the instruction looks at the value contained in the *demux_len* field 116 and uses this value referred to as *demux_val* as the 5 index into a table of $2^{\text{demux_len}}$ instructions. The instruction pointed to by the table pointer and the *demux_val* 118 is executed next. The use of a single instruction with an ASIC as described above lends itself to a convenient implementation of a graph-based filtering algorithm.

10. [037] Having described the methods of IP packet filtering according to the invention and its use in a system that utilizes a single instruction to perform expression matching, reference is made to Figure 5 which is a process flow diagram for the method of building a decision tree according to the present invention. The method, denoted generally as 150, begins at step 152 wherein the decision structure of the tree-like architecture used in the 15 decision tree is devised. Next, process flows to step 154, wherein a byte is chosen to subdivide in an effort to narrow down the filter set to only those filters that match. This is done by choosing the next byte to subdivide dynamically, and at each stage, based on the minimum wildcard heuristic. In particular, at step 156, the *cmpdemux* for all filters of this 20 stage is generated and, based upon the choice of byte, a calculation as to the maximum number of leftmost bits which all of the filters have in common in the byte is made. This calculation becomes our *cmp_len* value referred to in Figure 4.

25. [038] Next, at step 158, the longest mask length is determined among the current filters and the variable *demux_len* is calculated, step 160. As discussed above, *demux_len* equals *max_mask_len* minus *cmp_len*. At this point, all filters can be grouped into $2^{\text{demux_len}}$ sets, step 162, such that each *demux_val* maps to the appropriate set of matching filters for the chosen byte.

[039] Process flow is directed to step 163 wherein the decision is made if all filters sets have been processed over all bytes being filtered. If not, process flow is redirected back to step 154 and the algorithm recursively processes each of the remaining sets until all 5 sets have been processed over all bytes. Once all sets have been processed over all bytes, process flow is directed to step 164, wherein all subdivisions for all filter sets have been completed.

[040] Since wildcards are a very common occurrence within differentiated services filters, 10 it is not necessary to include wildcards in the calculation of the *cmpdemux* parameters (*cmp_len*, *demux_len*, etc.). Instead, the parameters can be calculated without the wildcarded filters and the filters added to each of the subdivisions with an additional subdivision made which represents the no-match condition of the *cmpdemux* instruction. 15 This is appropriate since the incoming data packet will automatically match any wildcarded filters. This special treatment of wildcards prevents all of the *cmpdemux* instructions from using a *demux_len* of 8, yielding a *demux* table of 256 entries.

[041] Figure 6 illustrates the use of the *cmpdemux* instruction in an example using filters 20 *f*1, *f*2, and *f*3 to prune an incoming data packet. The first *cmpdemux* instruction looks at byte 1. Since all filters contain a value 1 in this byte and a mask of $(11111111)_2$, then we would set the *cmp_len* to 8, because all of the filters contain 8 bits in common in this byte. 25 However, since the *cmpdemux* instruction's *demux_len* must be at least 1, *cmp_len* is set to 7 and the *demux_len* to 1, leading to a table of two entries. Since all filters fall into the category of having a *demux_val* of 1, it is specified in table entry 1 that the pruned filter set is {*f*1, *f*2, *f*3} (i.e., no filters were pruned). Since no filters have a *demux_val* of 0, that table entry corresponds to no filter matches.

5 [042] When a packet arrives, the *cmpdemux* instruction will look at the first *cmp_len* (i.e., 7) bits of the first byte of the search key and check to see if the value is *cmp_val* (0). If so, the instruction looks at the next *demux_len* (1) bits and uses that value to jump into the table to find the next instruction. So, if this packet does indeed contain a 1 in this byte, the next instruction is executed, labeled as 2 in figure 6. However, if the previous comparison fails, then we jump to the *NoMatchCase*, which specifies that no filters match this packet.

10 [043] The second *cmpdemux* instruction is constructed by looking at the values of the second byte in the current filter set, $\{f_1, f_2, f_3\}$. These filters have values $2 = \underline{0000} \underline{0010}_2$, and $5 = (\underline{0000} \underline{0101})_2$ in this position. These values have the first 5 bits in common, hence $15 \quad \text{cmp_len}$ is set to 5, cmp_val to 0 and demux_len to 3. Now the table contains 8 entries, with the 010 entry representing a match of $\{f_1, f_3\}$, the 101 entry representing a match of $\{f_2\}$, and all other entries corresponding to no matches to any of the filters. Since the 101 entry only contains f_2 , and since f_2 has a mask length of 2 bytes, processing is done and a leaf of the graph has been reached. In the case of the 010 entry, the $\{f_1, f_3\}$ filter set is further subdivided since one of the member filters (f_1) contains more bytes to match. In this case, the *cmpdemux* instruction labeled by 3 is executed.

20 [044] In the *cmpdemux* labeled 3, we look at values of the members of the current filter set, $\{f_1, f_3\}$, in byte 3. Since f_1 contains $3 = (0000011)_2$ in this byte and f_3 contains a wildcard (0/0), then it is sufficient to set up the *cmpdemux* to look at the first $\text{cmp_len} = 7$ bits of the incoming packet's value in the corresponding byte and see if it contains $\text{cmp_val} = 1$. Hence the demux_len is 1. Entry 0 in the table corresponds to matching filter f_3 . Since f_3 has no remaining byte to compare, processing is done. Entry 1 in the table corresponds to matching filters $\{f_1, f_3\}$. In this case, since there are more bytes to process in f_1 , the next *cmpdemux* instruction is executed, labeled as 4. Notice that f_3 is a wildcarded filter for byte 3 and was replicated in all table entries for this *cmpdemux*. In the

case that the incoming packet does not match *cmp_val*=1 for *cmp_len*=7 bits, the process jumps to the *NoMatchCase* which corresponds to matching filter f3.

5 [045] Finally, in the *cmpdemux* labeled 4, *cmp_len* is set to 7, *cmp_val* to 2, and *demux_len* to 1, similar to step 3. In this case, all table entries represent leaves in the graph and specify final filter match sets.

10: [046] It should be understood that tree packet filtering methods can be readily implemented in a hardware system, such as a router, to provide a data packet classification apparatus for use in applications such as Differentiated Services. With reference now to Figure 7, therein is shown a general block diagram for a router capable of use in a data network that allows the implementation of differentiated services by filtering incoming data packets. Specifically, the router 200 includes an interface 210 to data links 230 coupling the router 200 to other parts of the data network 250. Interface 15: 210 provides a point of ingress to incoming data packets. Router 200 is shown to include a processing means in the form of a packet processing unit (PPU) 212 which can be an application specific integrated circuit (ASIC) or general purpose computer processing unit (CPU) adapted to perform the IP packet filtering functions as described herein.

20 [047] Typically, a set of software instructions in the form of logic 214 is provided to operate in conjunction with ASIC 212 to carry out such functions. Logic means 214 can utilize filters 216 in order to construct a decision tree having a tree-like architecture such as tree-like architectures 50, 100. Of course the specific structure and form of the tree-like architecture would depend on the type of expression matching which is required based on, 25 for example, the level and/or type of service differentiation needed. Therefore, logic means 214, PPU 212, and filters 216 can be used to compare segments (or bytes) of incoming data packets to other values stored in a tree-like hierarchy such that the

segments are compared along stages of the tree-like hierarchy and not along every possible edge of a tree-like hierarchy to allow efficient and speedy filtering of incoming data packets. After matching and differentiation, PPU 212 can transmit data packets through egress port 220 which may involve transmission of a data packet to a specified
5 location based on its classification.

[048] The objective of the classification is to determine which of the filters match the current input classification key. As stated above, a classification key is composed of packet fields located at pre-specified offsets within the packet. One skilled in the art can use the methods of the present invention to create a classification key system that tests each byte of the key, in the order determined by the decision tree. The code to perform such a comparison consists of cascaded *cmpdemux* instructions, each of which narrows down the matching filter list. As discussed above, the *cmpdemux* instruction looks at the first *cmp_len* bits of a byte and checks to see if they equal the *cmp_val* quantity. If they
10 match, then take the next *demux_len* bits and use them as an offset into the *demux* table, pointed to by the *cmpdemux* instruction, to obtain the next instruction to execute. If the
15 *cmp_len/cmp_val* comparison fails, then we jump to the location specified in the *Cmp_fail_ptr*. The number of *cmpdemux* instructions that any particular code path will encounter is bounded by the number of bytes in the input key since each *cmpdemux*
20 divides the filter set according to a particular byte and no further comparison using that byte is necessary.

[049] The code block can contain a simple NoMatch handler, that is, the code that gets executed when a *cmpdemux* instruction fails. This handler is merely a jump instruction
25 which will jump to the instruction pointed to by the register referred to by the *CMP_FAIL* macro in the code. Thus, prior to each *cmpdemux* instruction, the *Cmp_fail_ptr* is loaded with the appropriate NoMatch case address.

[050] One possible example related to the packet filtering problem would occur if a particular set of fields (e.g., source or destination IP address) are possibly distinct in a filter set, but always fully masked (e.g., never containing any wildcard bits). In that case the 5 minimum wildcard heuristic will automatically select a byte from (one of) the non-wildcared field(s) as the first byte to search. The execution time of the search graph generation across the entire filter set may be avoided (at the possible expense of a minor increase in the size of the total search graph) if the heuristic is disabled for bytes within the non-wildcared field(s), and the bytes are searched in some fixed order.

10
15
20
25

[051] If the size of the field is reasonably large (e.g., ≥ 2 bytes), but the set of possible values in the filter set is small relative to the size of the set, then a more compact search method for that field could be employed. For example, a hash of the field could be used as the top-level branch, followed by one or more multi-byte comparisons (the number depending on the hash collision probability). This is an effective way of flattening the height of the search graph by using wider comparisons.

[052] The algorithm described above does not store intermediate results of the number of filters entirely matched, but instead maintains a logical association at each node with 20 the number of feasible filters (those that are not excluded from a match based on the previous comparisons). This approach is not efficient if the filter set has a pathological structure. For example, assume the classification key consists of two elements (e.g., packet header fields). A filter set of N filters consists of N_1 filters specifying a value for the first element, and wildcard in the second, while a set of $N - N_1$ filters specifies a value for 25 the second element, and is wildcared in the first. The set pruning graph algorithm will necessarily generate a graph with $N_1 * (N - N_1)$ leaf nodes, each corresponding to a

possible match of one or two filters (each from a different key set). For large numbers of N this can become prohibitive in memory and graph computation time.

5 [053] This particular problem can be reduced by storing intermediate results. In the example above, since the two sets of filters have disjoint masks, the classification problem could be broken into two independent search operations each on a separate classification key element. Then the results of the two operations could be combined, generating the filter match set.

10 [054] In a more general case, a filter set may satisfy the property that no individual byte in the set is wildcarded in every filter, but may be wildcarded in a majority of them. The classification problem could be simplified by splitting the filters into disjoint sets, based on the minimum wildcard heuristic, and performing two or more independent search operations. Alternatively, the splitting of the sets could be performed at each level of the graph. This might allow for the search graph of a set of bytes (wildcarded in the previous 15 searches) to be collapsed into a new graph, possibly testing one or more bytes multiple times (trading search graph depth for memory consumption). At the collapse point the set of fully matched filters is saved, and is later concatenated with the next match set. This process can be repeated multiple times as needed, based on the tradeoffs of execution 20 speed and memory needed.

25 [055] While the invention is described in connection with specific embodiments and with reference to specific diagrams and examples, it should be readily understood that changes and variations in the invention can be devised by those of ordinary skill in the art. For example, while the invention has been described in connection with methods for filtering incoming data packets by searching and comparing against "bytes" of the packets, it should be readily understood that an algorithm or method according to the invention may

search and compare based on fields smaller than or larger than a byte. Other variations will be apparent to those of ordinary skill in the art. Therefore, it is contemplated that such changes and variations will be within the spirit of the present invention and captured within the scope of the following claims.

5 0 0 0 1 2 0 7 0